IMPLEMENTING PREDICTIVE PROCESSING AND ACTIVE INFERENCE: PRELIMINARY STEPS AND RESULTS

Beren Millidge
Department of Informatics
University of Edinburgh
United Kingdom

March 11, 2019

ABSTRACT

Initial and preliminary implementations of predictive processing and active inference models are presented. These include the baseline hierarchical predictive coding models of (Friston, 2003, 2005), and dynamical predictive coding models using generalised coordinates (Friston, 2008a,b; Friston et al., 2010; Buckley et al., 2017). Additionally, we re-implement and experiment with the active inference thermostat presented in (Buckley et al., 2017) and also implement an active inference agent with a hierarchical predictive coding perceptual model on the more challenging cart-pole task from OpanAI gym. We discuss the initial performance, capabilities, and limitations of these models in their preliminary stages and consider how they might be further scaled up to tackle more challenging tasks.

1 Introduction

Predictive processing is a paradigm that has arisen in the last two decades and combines threads of research from Bayesian statistics, machine learning, and neuroscience to try to present a unified perspective on brain function. It began with the work of Rao and Ballard on predictive coding models (Rao and Ballard, 1999) and was proposed as a unifying theory of cortical function in (Friston, 2003, 2005). These predictive coding models were later extended to dynamical models predicting generalised coordinates in (Friston, 2008a,b; Friston et al., 2010, 2007) as a form of generalized filtering. The paradigm was then further extended to action in the form of active inference (Friston et al., 2011, 2009; Brown et al., 2011; Pezzulo et al., 2015; Friston et al., 2017) where it claims to be able to explain and model epistemic rewards and purposeful exploratory behaviour (Friston et al., 2015, 2016, 2017), as well as various aspects of choice behaviour (Friston et al., 2013). The core underlying concept of this theory is the free-energy principle, which states that every complex system must minimize a quantity called its variational free-energy in order to maintain homeostasis with the environment. This variational free energy is a tractable upper bound on the entropy of the system and thus by minimizing the free-energy the system can minimize its entropy and thus keep the boundaries of its system within homeostatic bounds (Friston and Ao, 2012; Karl, 2012; Friston, 2010; Friston and Stephan, 2007; Friston, 2009). Despite the impressive mathematical and theoretical grounding of the theory, and the extent of its neuroscientific claims there have been relatively few concrete implementations of the models proposed in the literature. Some exceptions are the implementations of dynamical and active inference agents proposed in (Baltieri and Buckley, 2017; McGregor et al., 2015; Buckley et al., 2017), predictive coding models including relatively large scale ones implemented by Spratling (Spratling, 2008, 2017), and the the SPM repository – https://www.fil.ion.ucl.ac.uk/spm – which contains

implementations of most of the proposed models in the literature. Even where the models have been implemented they have not typically been applied to complex large scale tasks of the type that is regularly tackled by machine learning and related areas. This is a big gap since, if the theoretical claims of predictive processing are true and these theories and models are in some sense implemented in the brain, then they should be able to scale up to solve large-scale complex challenges similar to those currently solved in machine learning.

In this review and in other work we aim to take a step towards closing this gap. We present preliminary results from our own implementations of a variety of predictive processing models in the literature. These include hierarchical predictive coding models (Friston, 2003, 2005; Rao and Ballard, 1999), dynamical predictive coding models using generalized coordinates (Friston, 2008a; Friston et al., 2010) and also a few forays into active inference.

This work is still at a very early stage, however preliminary results are promising. In general, the results show that even in the simplest stages the models do work and can learn to represent and predict stimuli correctly. The prediction errors do converge to zero, indicating strongly that the learning rules are correct and are practical to use even in high dimensional cases. Performance is far from perfect however, and the models are typically very sensitive to a range of initial conditions, parameter initializations, and hyperparameter settings. In general the models are usually exceptionally sensitive to the precision parameters, and often when the precisions themselves are updated using the gradient descent learning rules proposed the models fail to converge. There still needs to be substantial work in scaling up these models, and making them more robust and performant before they can even become remotely competitive with current state-of-the-art methods in machine learning. Nevertheless, we believe that the implementation and improvement of predictive processing models and theory is a valuable and promising research direction, and that these methods have several advantages over current machine learning methods. Predictive coding perceptual models are unsupervised and biologically plausible, can be optimized entirely in an online and layer-wise fashion well-suited to parallelization, are generative, and can easily represent temporal dimensions of stimuli in terms of generalized coordinates. Active inference approaches to action promise to solve the exploration/exploitation dilemma through a principled approach to valuing purely epistemic rewards, and also enable purposeful exploratory behaviour. They also enable flexible specification of behaviour and goals which can be switched out "at runtime" and then inferred dynamically as opposed to current reinforcement learning methods which typically rely on learning a fixed policy for a given reward signal which cannot easily be adjusted or swapped out without substantial retraining. Finally, these methods hold out the promise of being able to implement computational algorithms which not only could be implemented in the brain,but which have been proposed as an explanation of a substantial part of cortical function, and thus could be a substantial step towards constructing general computational intelligences.

This report comprises three sections. The first presents results with static hierarchical predictive coding models. These simply try to predict the next input given the current input. They are tested on relatively high dimensional data in the form of images from the MNIST and CIFAR10 datasets. We show that they are able to learn to predict and represent images, effectively acting akin to autoencoders in machine learning, but with solely local and biologically plausible predictive coding learning rules. They are also able to learn effective latent representations of their input and can be used to generate sample digits and images which they have not been presented with. The second section presents results with dynamical predictive coding models. These models represent the input in generalised coordinates - i.e. as vectors of all the temporal derivatives of the input. Current results are preliminary and concerned with predicting sine waves and other simple wave-forms for which analytic temporal derivatives are known. The models show some success at rapidly learning to predict these simple waveforms. We also implement "full-construct" models (terminology taken from Buckley et al. (2017), which are predictive coding models that are both dynamic and hierarcical), and preliminary results indicate that they are also able to sucessfully learn to predict simple waveforms. The third section presents initial results for active inference. Active inference is a paradigm which extends predictive processing and the free energy principle to action. In active inference there is an actual agent that interacts directly with an environment in such a way as to minimize its free energy through both perception and action. We implement and discuss two active inference models. The first is a re-implementation of the model given in Buckley et al. (2017), of an active inference

thermostat which seeks out a given temperature on a lineworld of decreasing temperature. The second, and significantly more complex model, is one that solves the cart-pole environment of OpenAI gym, a baseline reinforcement learning environment. It uses a hierarchical predictive coding generative perceptual model to predict its observations over future timesteps and uses these predictions to select actions which minimise its variational free energy. We show that it performs reasonably well on the cart-pole task and that these results indicate that active inference shows promise as a method that could be scaled up to perform well on challenging reinforcement learning tasks.

Results presented in this report are preliminary, and much work remains to be done in establishing the precise capabilities and limitations of the models presented, how to scale them up to attack larger tasks, and how to improve their robustness and performance.

2 Static models

2.1 Introduction

Static models are the simplest predictive processing models. They try to predict the incoming sensory inputs at time t+1, compare their predictions with the actual input and then adjust their parameters - weights and 'cause units' so as to minimize this prediction error. The model can be made hierarchical such that there are multiple level of 'cause-units' involved in the prediction. In hierarchical models the prediction errors are passed upward to the level above while predictions are passed downwards. Each layer tries to predict the input to the level below, and adjusts its own parameters to minimize the prediction error. These models were introduced first in Friston (2003) with a thorough review in Friston (2005). They generalise predictive coding as proposed in Rao and Ballard (1999). These models can be derived from the general form of the variational free energy through gaussian assumptions about the variational density Q and then applying a Laplace approximation to the gaussian generative density (Friston et al., 2007).

Each layer of the static hierarchical predictive coding model is composed of 'cause units', denoted by μ which are meant to represent the latent causes of sense data in the environment and weights, denoted θ. Predictions are denoted with a hat – the incoming sense data is denoted ϕ and the predicted sense data is denoted $\hat{\phi}$.

Each prediction is defined to be some function of a linear combination of the weights and the cause units: $\hat{\phi} = f(\theta.T\mu)$. This is extremely similar to the standard rule used in neural networks for computing the activations of a layer.

These layers can be stacked hierarchically, such that the cause units of one layer are the predictions of the layer above. This can be expressed mathematically as:

$$\hat{\phi} = f(\theta_1.T\mu_1) \tag{1}$$
$$\mu_1 = f(\theta_2.T\mu_2) \tag{2}$$
$$\mu_2 = f(\theta_3.T\mu_3) \tag{3}$$
$$... \tag{4}$$

In the linear case the function f is just the identity so that the prediction is simply a linear combination of the cause units.

The prediction errors for a level are defined as the difference between the activation for the level and the predicted activation. For the first level of the hierarchy, the prediction error is the difference between the predicted and actually received sense-data.

$$\epsilon = \phi - f(\theta^T \mu)$$

The variational free energy of such a model is a simple quadratic sum of prediction errors, divided by their precision. For a full derivation see (Friston, 2005, 2003; Buckley et al., 2017; Bogacz, 2017).

$$VFE = \sum_i \epsilon_i^T \Sigma_i^{-1} \epsilon_i$$

Where the ϵ_i is the prediction error of the i'th layer of the network and Σ_i^{-1} is the precision, or inverse-variance of that layer of the network. This variational free energy can be minimised by standard gradient descent, and derivatives can be quiet easily found for the parameters μ - or the 'cause-units':

$$\frac{dF}{d\mu_i} = \Sigma_i^{i-1} \epsilon_i^T \frac{df}{d\mu} \theta_i + \Sigma_{i+1}^{-1} \epsilon_{i+1}$$

And for the θ parameter, or the 'weights' the gradient update is:

$$\frac{dF}{d\theta_i} = \Sigma_i^{-1} \epsilon_i \frac{df}{d\theta_i} \mu_i^T$$

One can also derive a gradient update for the precisions:

$$\frac{dF}{d\Sigma_i} = -\Sigma_i^{-T} \epsilon_i \epsilon_i^T \Sigma_i^{-T}$$

These equations provide the simple update rules for the static model. For a full derivation of these equations, see (Friston, 2003).

Hierarchical predictive coding models were implemented which learned using these equations. They were tested on the image datasets MNIST (a dataset of 60,000 handwritten digits - 28x28 pixels in size) and CIFAR10 (an image dataset of ten classes of everyday objects which are 32x32 pixels in size).

2.2 MNIST

Hierarchical static predictive coding models were trained on MNIST digits. Since there was no temporal component to the images, the predictive processing objective effectively reduced to that of a simple autoencoder except that the model learned in an unsupervised manner using the equations above instead of through stochastic gradient descent.

The results below were generated with a two layer linear static predictive processing network with the first layer having a dimensionality of 784 (the dimensionality of a flattened 28x28 MNIST digit) and the second layer having a dimensionality of 20. Thus the predictive processing model learned to map the 784 dimensional vector reprsenting the pixels of the image into a 20 dimensional latent space and then use that to predict the digit. The network was trained on 10000 digits and ran for 100 epochs. The learning rate was set to 0.01.

The model was able to recreate MNIST digits successfully, as shown in the example reconstructions below:
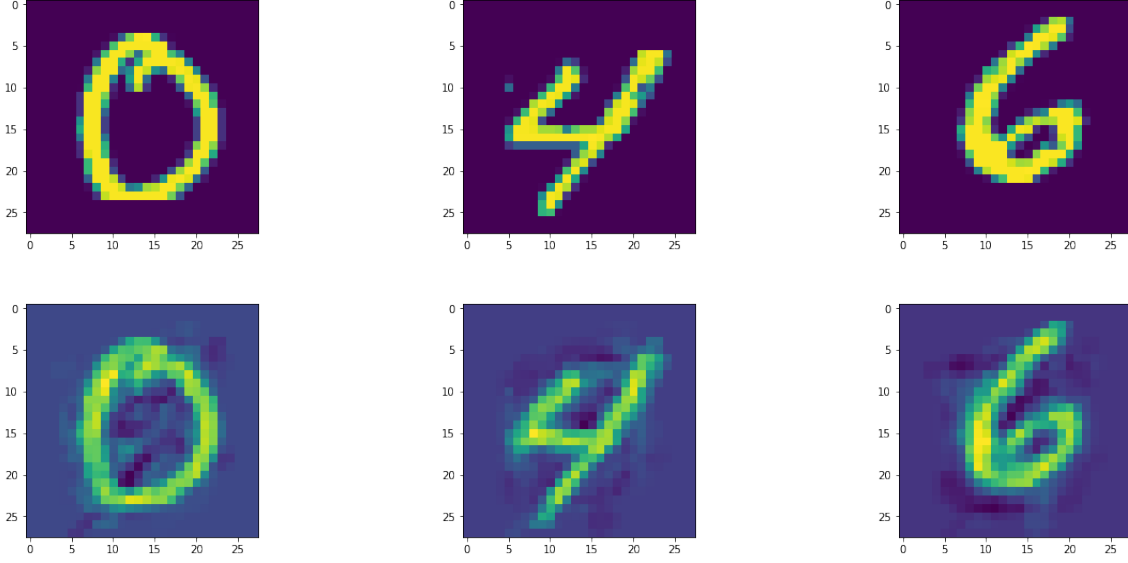
4

Figure 1: MNIST digits in the training set recreated by the network. Top row the actual digits, bottom row, the predictive reconstructions.

Additionally, the model was also able to recognize and reconstruct previously unseen MNIST digits, albeit with slightly lower fidelity. Nevertheless it is impressive how rapidly and well the network is able to generalize to completely unseen digits.
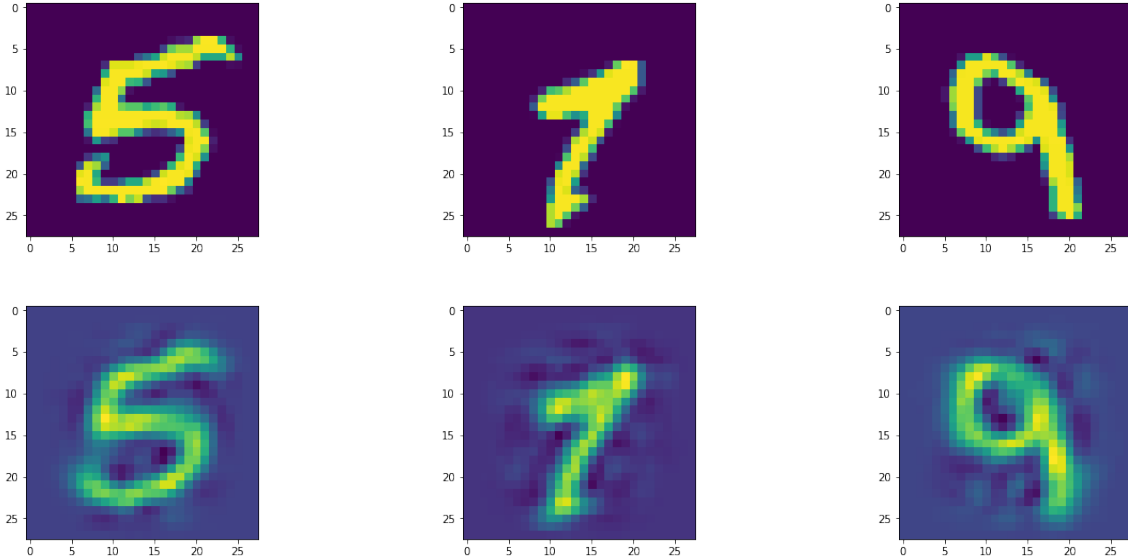

Figure 2: Unseen MNIST digits in the test set recreated by the network. Top row the actual digits, bottom row, the predictive reconstructions.

Since the model is a generative model, it is also able to generalize outside the training set to generate, or 'dream', completely unseen digits by sampling from the latent space. Examples are shown below:
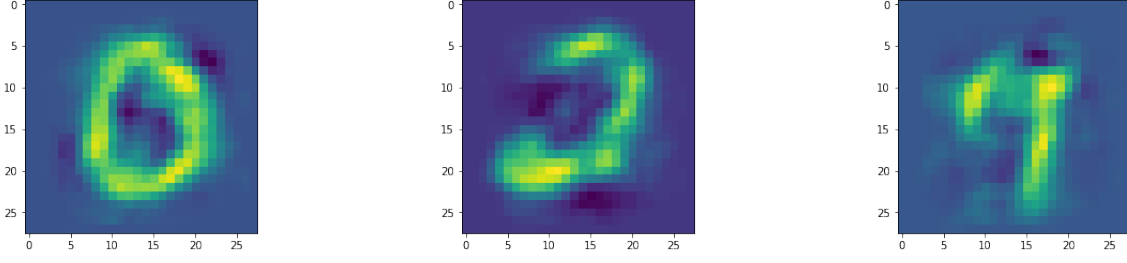
Figure 3: Images of hallucinated digits "dreamt" by the network. These were generated by sampling the latent space around the representations of some exemplar digits in the latent space.

It is also possible to interpolate between digits in the latent space. To do this, the latent representations of the two images are taken and a point in the latent space is picked which is slowly updated from the initial image towards the final image. At each step in the process, the point in the latent space is propagated forwards to produce the reconstructed image expected from that point in the latent space. In the supplementary materials there is a video of such an interpolation, where the network interpolates between a 3 and a 9.

Experiments were also conducted with larger and nonlinear networks comprising sigmoid nonlinearities. The results are largely similar. Better results are obtained with larger and deeper networks but training times are also significantly longer. Because of this, relatively little exploration of the capabilities of deep nonlinear hierarchical predictive coding models was undertaken. This is a rich area for future research, especially given the vast increases in expressiveness that come with depth and nonlinearity in artificial neural networks.

Finally, to visualize the latent space, PCA (principal components analysis) was applied to the learned representations in the 20 dimensional latent space to shrink it down to a two dimensional space. It is apparent upon inspection of figure 4 that the latent space, even when shrunk down to two dimensions, does a good job of clustering the MNIST, putting all the 1s in the top left corner, or all the zeros in the middle right. This strongly indicates that the predictive coding model is able to learn the categories of digits despite being trained in an entirely unsupervised way without any knowledge of the true identities of the digits.
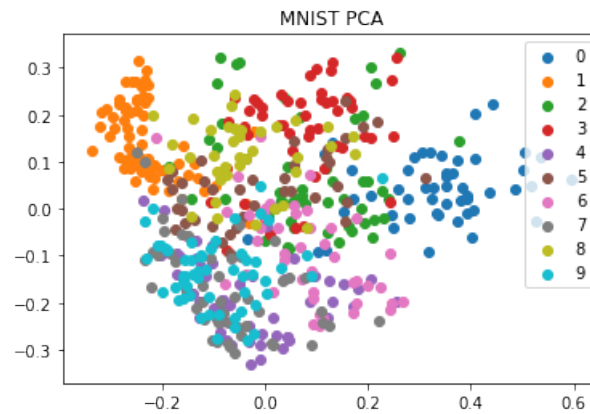

Figure 4: A PCA clustering plot of the values of test MNIST digits in the latent space. Even though the 20 dimensional latent space has been reduced down to two, clusters are still visible. For instance, all the 1s are clustered in the top left corner.

6

2.3 CIFAR

Similar experiments were conducted on the CIFAR10 dataset with predictive coding models extremely similar to that used on MNIST. We only experimented with linear predictive processing models for CIFAR. The results are as follows:

The predictive coding model used in the CIFAR experiments was comprised of two predictive coding linear layers with dimensionality 1024 (32x32) in the first layer, and with a latent space of dimension 50. The network was trained on 1000 CIFAR10 images for 100 epochs. The learning rate was set to 0.01.

The hierarchical predictive coding CIFAR models are able to learn to reconstruct CIFAR images with impressive fidelity given that they were compressed from a 1024 dimensional image into a 50 dimensional latent state.
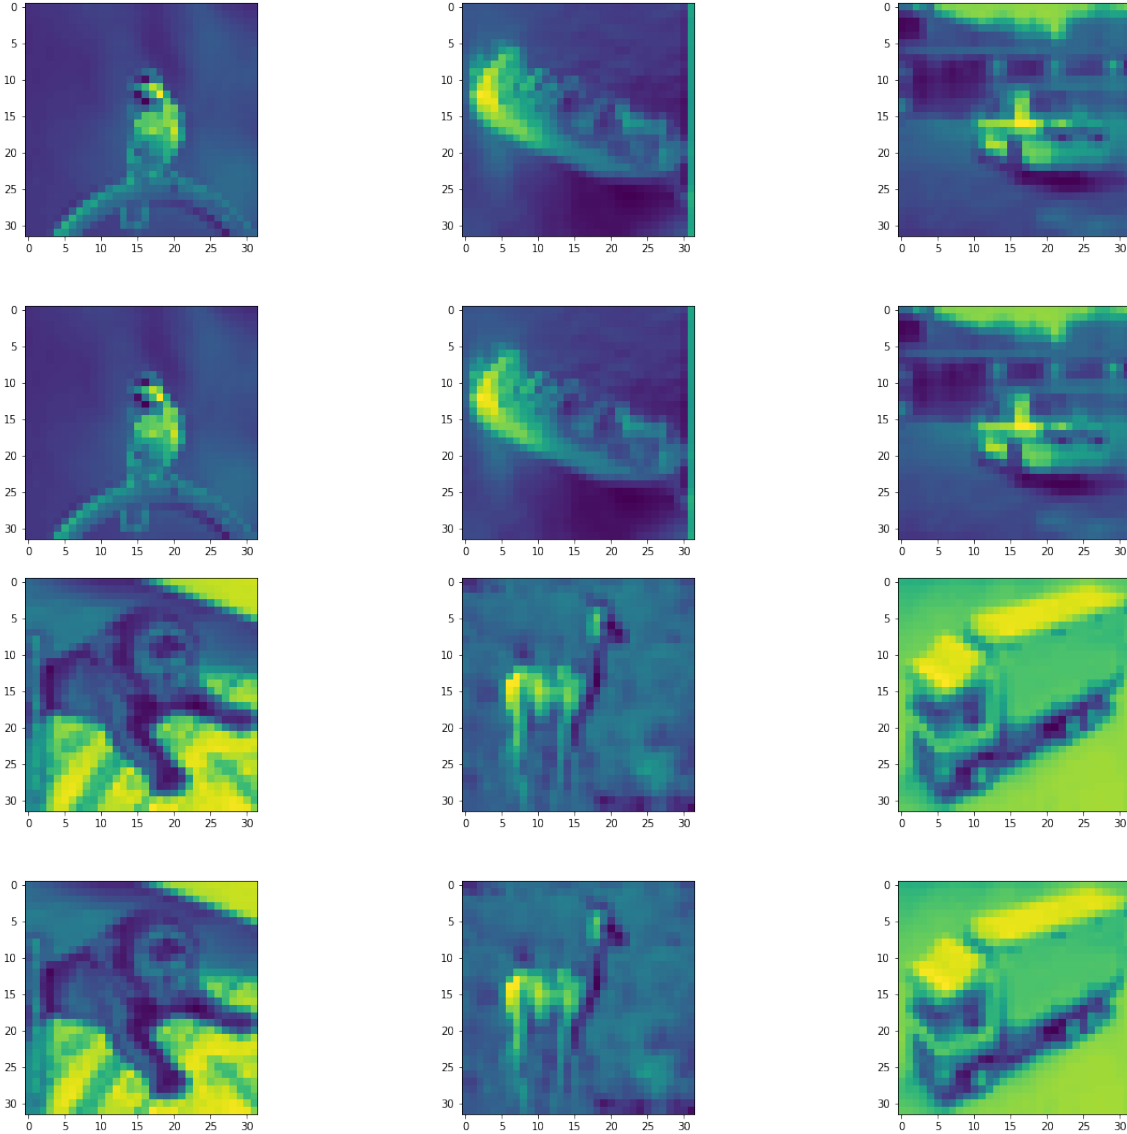

Figure 5: Test set CIFAR digits reconstructed by the network. The first of the two lines is the image and the second is the reconstruction. The network is extremely good at reconstructing CIFAR images.

As with MNIST, the CIFAR models are able to interpolate in the latent space. An example interpolation is shown below, in figure 6, where the model is interpolating between a horse and cat.

Figure 6: The Cifar model interpolating between a horse and a cat. Read the images left to right top to bottom - like text. The interpolation is done by stepping in the latent space from the representation of the first image in the direction of the second until it is reached.

Overall, preliminary implementations of linear and nonlinear hierarchical static predictive coding models have been completed and tested. Initial results are encouraging. The models are able to learn robust and generalizable latent-space encodings of their inputs using the unsupervised predictive coding update equations. In the case of MNIST, we have shown that such models are able to learn a clustering by digit identity in the latent space, despite being trained in an unsupervised manner without such information. The models are also able to generate good reconstructions, including of unseen images drawn from the test set, to interpolate between images, and also to generate new unseen digits by sampling from the latent space of their representations. The unsupervised predictive coding objective does appear to give good learning performance, and typically converges rapidly. Much future work remains to be done in scaling these networks up in terms of size, and depth, and including nonlinear activation functions, and tuning them for performance and speed. Nevertheless, initial results seem promising.

3 Dynamical Models

3.1 Introduction

Dynamical predictive coding models predict not only their inputs at time $t + 1$ but also all higher temporal derivatives of their inputs. This means that they are sensitive to and track rates of change of the incoming data (and rates of rates of change, and so forth). They thus implement generalized filtering as described in Friston (2008a,b); Friston et al. (2010).

Dynamical models can also be hierarchical, which we refer to as "full-construct" models following the excellent mathematical review of Buckley et al. (2017). The hierarchical dynamical model tries to predict the dynamics of not only the sensory causes but those of its own predictions at progressively higher levels of abstraction. At each level the prediction errors of that level, and thus the variational free energy, is minimized.

The key concept is that of generalized coordinates. Generalized coordinates extend the notion of the sense-data to not only include the sense-data itself, but all temporal derivatives of the sense data. The input data then becomes an infinite dimensional vector comprised of the data and all the temporal derivatives, which will be denoted as $\tilde{\phi}$. Where:

$$\tilde{\phi} = [\phi, \frac{d\phi}{dt}, \frac{d^2\phi}{dt^2}, \frac{d^3\phi}{dt^3}....]$$

To account for this, the cause units and weights also need to become infinite dimensional vectors of their temporal derivatives - i.e.

$$\tilde{\mu} = [\mu, \frac{d\mu}{dt}, \frac{d^2\mu}{dt^2}, \frac{d^3\mu}{dt^3}....] \tag{5}$$

$$\tilde{\theta} = [\theta, \frac{d\theta}{dt}, \frac{d^2\theta}{dt^2}, \frac{d^3\theta}{dt^3}....] \tag{6}$$

The aim and general construction of the predictive processing model is still the same - to try to predict the incoming sensory data $\tilde{\phi}$ with the prediction.

$$\hat{\tilde{\phi}} = \tilde{f}(\tilde{\theta}\tilde{\mu})$$

And then to minimize the dynamical prediction error.

$$\tilde{\epsilon} = \tilde{\phi} - \hat{\tilde{\phi}}$$

This prediction error can then be minimized at each dynamical level by the same mechanism and update rules as with the standard static predictive processing model. Things get more complex when we consider a model which has both dynamical and hierarchical components, which may have interactions between them. This we call a full-construct model following the lead of this tutorial (Buckley et al., 2017). In a full construct model there is a dynamical hierarchy of levels where it is assumed that each dynamical order is only able to affect the level below:

$$\phi = f(\mu; \theta) \tag{7}$$

$$\mu = f(\mu'; \theta') \tag{8}$$

$$\mu' = f(\mu''; \theta'') \tag{9}$$

$$... \tag{10}$$

Similarly, there is an equal hierarchical hierarchy, where each level is assumed to be predicted by the level above it:

$$\phi = g(\mu_0; \theta_0) \tag{11}$$

$$\mu_0 = g(\mu_1; \theta_1) \tag{12}$$

$$\mu_1 = g(\mu_2; \theta_2) \tag{13}$$

$$\mu2 = g(\mu_3; \theta_3) \tag{14}$$

$$... \tag{15}$$

Therefore, each node in the lattice of hierarchical and dynamic hierarchies is influenced by two separate predictions - the dynamical prediction going from higher dynamical orders to lower, and the hierarchical prediction propagating

9

from higher levels of the hierarchy to lower ones. Thus, a single state of a cause-unit μ_i^n, where i is the level of the hierarchy, and n is the dynamical order, is defined to be:

$$\mu_i^n = f(\mu_i^{n+1}; \theta_i^{n+1}) + g(\mu_{i+1}^n; \theta_{i+1}^n)$$

This means that the variational free energy must sum over both dynamical and hierarchical prediction errors, such that:

$$VFE = \sum_i \sum_n \epsilon_i^n \Sigma_i^{-1} \Sigma_n^{-1} \epsilon_i^n$$

And that additionally the updates for the representation-units and the weights must take this into account. The revised update rules are presented below:

$$\frac{d\mu_i^n}{dt} = \Sigma_{n+1}^{-1} \epsilon_i^{n+1} \frac{df}{d\mu_i^n} \theta_i^{n+1} + \Sigma_n^{-1} \epsilon_i^n + \Sigma_{i-1}^{-1} \epsilon_{i-1}^n \frac{dg}{d\mu_i^n} \theta_{i-1}^n + \Sigma_i^{-1} \epsilon_i^n \tag{16}$$
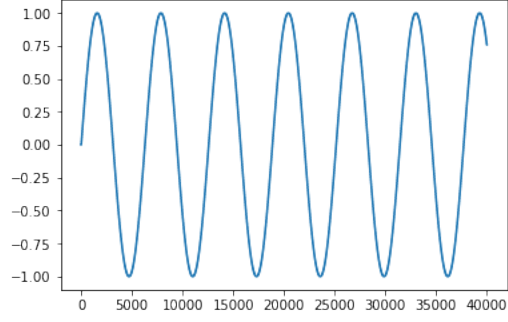
And for the weights the update rule thus becomes:

$$\frac{d\theta_i^n}{dt} = \Sigma_i^n \epsilon_i^n \left(\frac{df}{d\theta_i^n} \mu_i^{(n+1)T} + \frac{dg}{d\theta_i^n} \mu_{i+1}^{nT} \right)$$
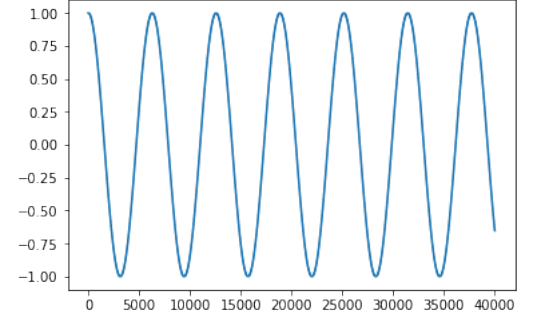
These rules appear somewhat more complicated than the corresponding rules in the static case. Nevertheless they are not particularly more difficult to compute.
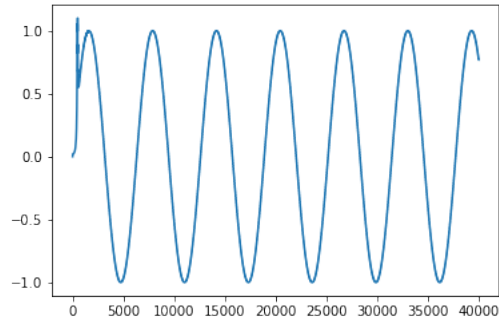
3.2 Predicting Sine waves

Preliminary dynamical and full construct models were implemented and tested on simple stimuli. The first task the dynamical models were tested on was predicting a sine wave. This is the perfect toy-task since sine waves have analytic derivatives to any infinite degree. A three layer linear dynamical model was implemented and learned to predict a sine wave and it's first two temporal derivatives. The model rapidly learned to predict the sine wave, as can be seen from the training graphs below. However, there was a consistent phase-error in the predictions it made, which could have been caused by the rapid rate of change of the sine-wave observations.

(a) Incomding sense data - i.e. a sine
wave

(b) First derivative of the incoming
sense data

(c) Predicted incoming sense data

(d) Prediction temporal derivative of the
incoming sense data

(e) Prediction error at the first dynamical
level

(f) Prediction error at the second dynam-
ical level

(g) Acitvation of representation units at
the first dynamical level

(h) Activation of the representation units
at the second dynamical level

Figure 7: Training graphs of the dynamical model on sine wave input

The dynamical model does not only work with sine waves. The model was also tested on more jerky waveforms sawtooth waves. In this case a two layer linear dynamical model was used which learned to predict the sawtooth wave and its first temporal derivative. The model predicts the wave very successfully, including the temporal derivative, although there it is a little less successful. Once again there is a persistent patterned prediction error, likely cause by a lag time between the models predictions and the uncertainty it receives. Solving this issue in dynamical models is still an open problem.
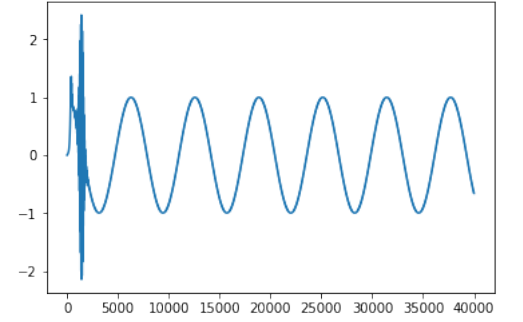

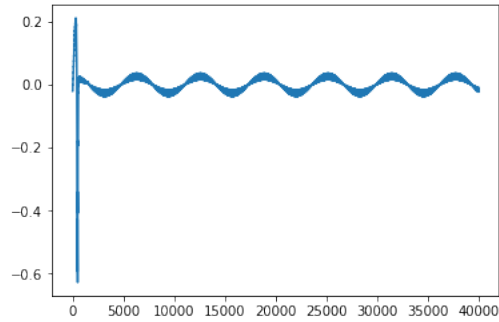(a) Incomding sense data - the sawtooth wave

(b) First derivative of the incoming sense data
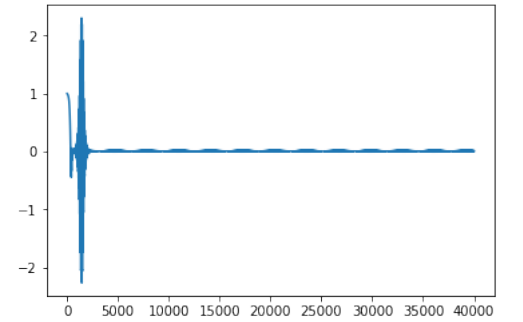

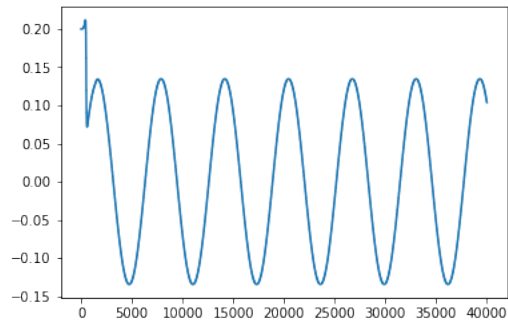(c) Predicted incoming sense data

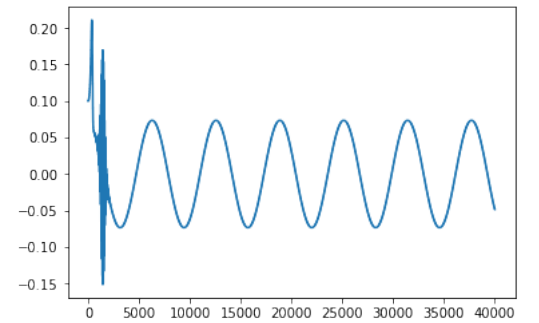(d) Prediction temporal derivative of the incoming sense data

(e) Prediction error at the first dynamical level

(f) Prediction error at the second dynamical level

(g) Activation of representation units at the first dynamical level

(h) Activation of the representation units at the second dynamical level

Figure 8: Training graphs for dynamical models

These graphs show that the model is able to predict the sawtooth wave very successfully, although it makes some initial mistakes.

In sum, the dynamical models do work, and can successfully represent both the sense-data and higher temporal derivatives of the data. They converge well and learn successful predictive models on toy tasks. Much work needs to be done however extending them to more realistic and challenging tasks.

3.3 Full Construct Models

Full construct models combine both hierarchical and dynamical models. They not only try to infer the dynamics of the sense data, but extend those predictions hierarchically, so that each level tries to predict the dynamics of the previous level up to any hierarchical level of abstraction. The model forms a lattice, with prediction errors being passed upwards both hierarchically and dynamically with predictions being fed downwards. Both dynamical and hierarchical predictions are integrated at each node. The system as a whole thus minimizes the prediction errors of both the hierarchical and dynamical models. The variational free energy of the model is comprised of a sum of both hierarchical and dynamical prediction errors, all weighted by their respective precisions.

Both linear and nonlinear full-construct models were implemented. The results below are reported only for linear full-construct models. This is because linear models trained faster and were also substantially more stable whereas nonlinear models tended to suffer from slow convergence or exploding gradients.

A full construct model was implemented and trained to predict sine waves. The model had two hierarchical layers and three dynamical layers. Training graphs are shown below:

(a) Incomding sense data - Sine wave

(b) First derivative of the sense-data

(c) Predicted incoming sense data

(d) The models' prediction of the incomign sense data

(e) The models' prediction of the first derivative of the incoming sense-data

(f) the models' prediction of the second derivative of the incoming sense-data

(g) The prediction error at the first hierarcical layer

(h) The prediction error at the second hierarcical layer

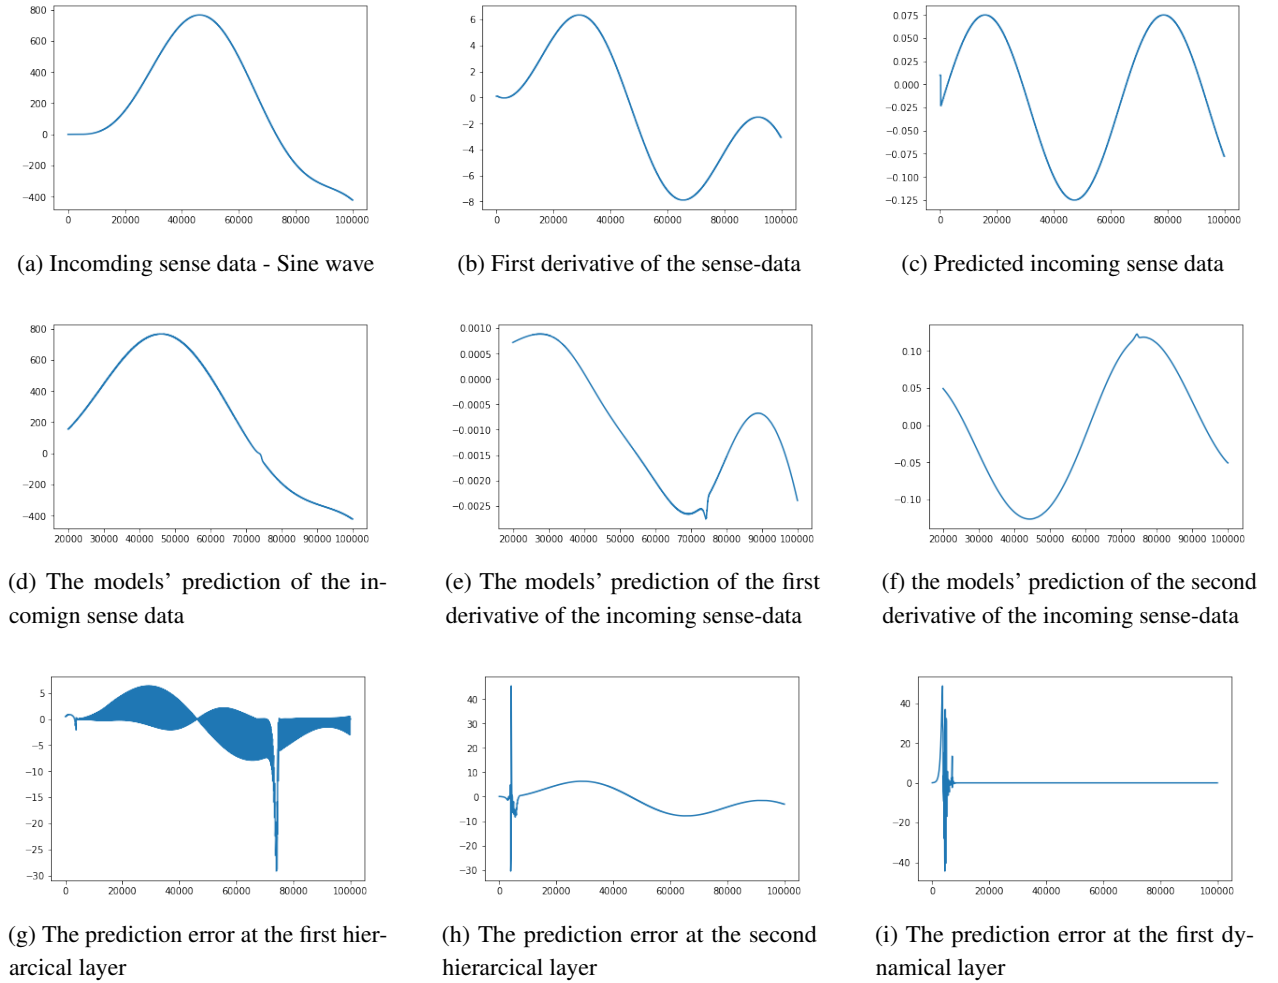(i) The prediction error at the first dynamical layer

Figure 9: The training graphs of the full construct model. It can successfully predict the first three temporal derivatives of a sine wave, and also minimise prediction error up to multiple hierarchical layers.

The top two rows of graphs show the incoming sense data and the first two temporal derivatives of the sense-data The next two rows show the prediction errors over time for various levels of the hierarchy. The full construct model appears a bit less stable and successful than the simple dynamical model, likely because it is much more complex and has many more moving parts. Nevertheless it manages to learn the sine wave shapes relatively faithfully and does also manage to rapidly reduce the prediction error over time. Moreover these sorts of tasks do not really play well to the strength of the full-construct model since the input data (the sine wave) contains no suitable hierarchical structure for the higher levels to model. It seems likely that as these models are scaled up to more challenging tasks, the greater expressivity and power of the full-construct models will become more apparent. Scaling up and making the full construct model more stable is an important area of future research. Nevertheless a major hurdle has been overcome by simply implementing it and getting it to converge in a simple prototype case.

4 Active Inference

Next we turn to active inference. Active inference extends the free energy principle to action. This means that not only do agents learn to perceive and organize their experience so as to minimize their variational free energy; they also act so as to sub-sample the sensory inputs they most expect, thus minimizing their free energy through action as well as perception. Active inference agents are thus in constant and active engagement with the world - emitting actions and receiving updated sense-data in return. The task faced by an active inference agent is quite difficult. It must create a predictive and generative model of the environment it is in, comprising one of the predictive processing models implemented previously. Moreover, since it can also act in the world it must, therefore, also have a generative model linking its actions to the state of the world, and then use this model to choose the actions that minimize the expected free energy.

Several different architectures have been experimented with, and there are many degrees of freedom in the design. Here we focus on two simple prototype active inference agents. The first is an re-implementation of a predictive processing agent presented in Buckley et al. (2017) which acts as a thermostat. It comprises a two level dynamical model which learns the sense-data and its first temporal derivative. It is provided a-priori with a precise action model - a model of how its actions affect the sensory data it will perceive in the future, and it can use this action model and its dynamical perceptual model to move to a region of desired temperature in a simple line-world.

The second agent was trained to complete a basic, but non-toy reinforcemnt learning envirnoment - that of the cart-pole environment from the OpenAI gym. It must learn from scratch both a perceptual model of how its observations evolve over time, but also a simple action model of how its actions at time t affect its observations at time $t + 1$. The architecture of this agent is a simple full-construct model with two dynamical and three hierarchical layers. The action model is implemented as a simple linear model feeding into the highest level of the full-construct model. It implements active inference by 'expecting' that the cart-pole is upright and stationary and then taking the action that will it infers will bring it closest to that reality.

4.1 Predictive Thermostat Agent

This agent is a re-implementation of the active inference thermostat presented in Buckley et al. (2017). The agent lives on a line world and has a position. Each position on the line has a temperature starting at 20 at position 0 and decreasing quadratically with distance from the beginning. The agent 'expects' to sense a certain temperature T and is able to act so as to update it's velocity in the world to reach its desired/expected temperature.

This task is a very simple one. It must move to a point where it can experience its expected temperature. The agent's state consists of two numbers - it's position and it's velocity, and it senses only two values - the temperature ,and the

rate of change of the temperature. It must choose a single value for the action - the velocity with which it will move for that timestep. It knows a-priori (i.e. the modellers provide it with) a correct model for how its actions interact with the world. The agent was extended beyond that presented in Buckley et al. (2017) to possess a full-construct dynamical model with 3 hierarchical and 2 dynamical layers. This means that the agent models both the temperature and the first derivative of the temperature with a 3-level hierarchical network.

This task can be completely solved with our active inference agent implementation. The agent starts at a position of 0 and tries to reach a desired temperature of 10 degrees. The training graphs of the agent are shown below.

(a) The temperature experienced by the thermostat agent.

(b) The X-position of the thermostat agent over time

(c) The action value taken by the thermostat agent

(d) The prediction error of the first dynamical layer

(e) The prediction error of the second dynamical layer

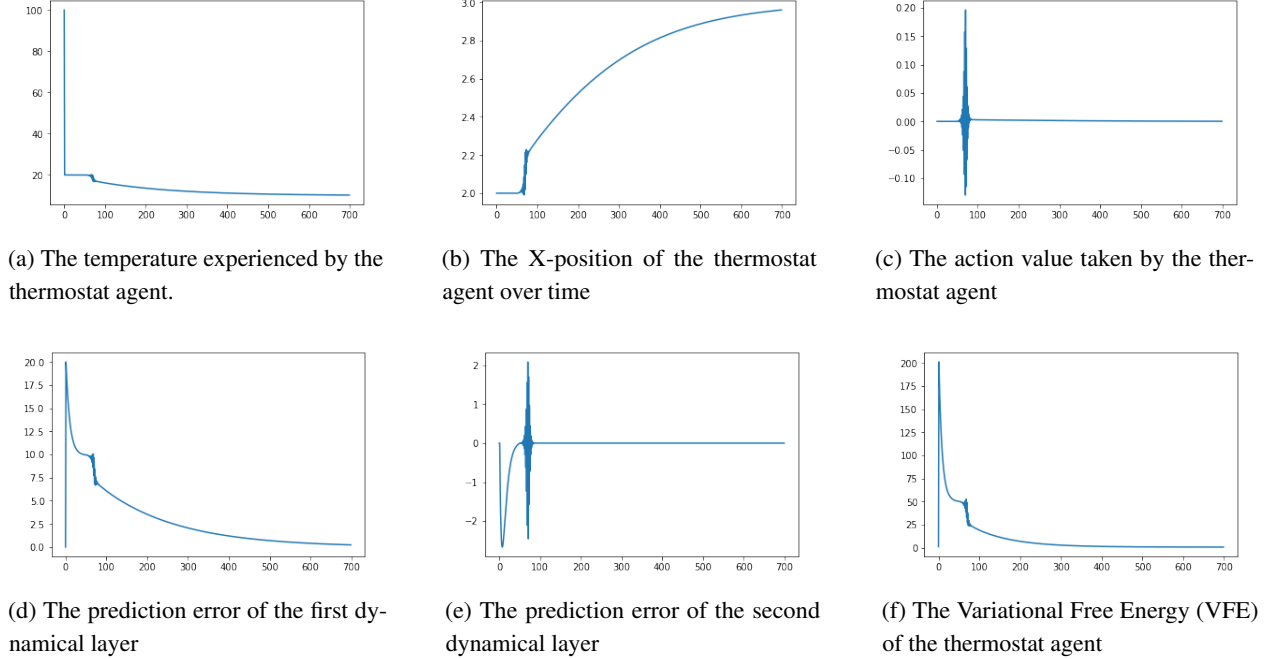(f) The Variational Free Energy (VFE) of the thermostat agent

Figure 10: The training graphs of the active inference thermostat model. The agent successfully asymptotically approaches the desired temperature of 10. The agent quickly learns to model its sensory inputs as can be seen by the rapidly decreasing prediction errors and VFE.

There are several interesting things to note in these graphs. The first and most important is that the agent is completely successful. It manages to asymptotically approach the correct temperature - in this case 10 degrees, but it can manage a wide range of desired temperatures. It also manages to successfully predict its environment, as shown by the graphs of the prediction errors which rapidly decrease to zero. The prediction error of the first dynamical layer only asymptotically decreases to zero since this is where the agents prior that it be at the desired temperature is fed into the system.

The agent can seemingly only converge asymptotically to the correct solution. This appears to be an artifact of the way active inference is setup here. Action is initiated through the prediction errors which given a good perceptual model, roughly corresponds to the difference between how the world actually is, and how the agent a-priori expects it to appear. Action is then proportional to the magnitude of the prediction error, at least in the case of linear predictive coding models. This naturally leads to an asymptotically decreasing prediction error, and thus an asymptotically decreasing action.

The algorithm is also somewhat sensitive to the learning rate. If the learning rate is too high, the weights and cause units are prone to explode. If they are too low then learning will proceed at an extremely slow pace and often fail to converge at all. The agent is also extremely sensitive to the precisions, especially the ratio between sensory and action precision. As an extreme case, there are two behaviours of the agent depending on this ratio - either the agent will converge to the desired temperature, or else the agent will forever move away from the origin, tending towards zero temperature. The transition between behaviours is especially abrupt. The action dynamical precision is set to 0.01. When the sensory precision is set to 0.70915 the temperature reached will converge to 0. When it is set to 0.70917 the temperature will converge to the correct value of ten. When it is 0.70916 it will be intermediate. It appears the switch in behaviour to the precisions is so sensitive to the precisions it happens in the fifth decimal point(!).

I suspect that the reason for such sensitivity to the precisions lies in a deep problem with active inference as currently formulated, which is that it intrinsically blends perception and action in a way that action can contaminate perception and vice-versa. When dealing with any prediction error the system can adjust the expectations of the world to match the prior or to act to change the state and thus the sense data to act the prior. The trouble comes when the system will try to both at the same time, which is what will inevitably happen in such systems when the precisions are not set so that one completely dominates the other. This means that the system minimizes prediction error through a compromise of action and sensory updating. The problem is that this compromise is usually deleterious for the agent. On the one hand, since the action will have been a compromise with the sensory updates, it won't be sufficient to actually fulfill the agent's goal but will always be slightly short of it. On the other, since the prediction error comes from an exogenous effect of desiring the world to be different to how it is, the perceptual update will almost certainly render the perceptual system less accurate as well. Thus the close integration of action and perception in this way can cause significant issues for the agent and its perception and action models.

This simple case has demonstrated an existence proof of an active inference agent which can solve a very simple task in a very simple world. It uses quite a complex full-construct hierarchical dynamic model which converges and accurately models the environment. It can then, by the process of active inference, act so as to minimize the prediction error, and if the precisions are set up correctly, it will then converge accurately on the final desired temperature.

4.2 CartPole Agent

After the existence proof of the thermostat agent, we extended the active inference paradigm to a more difficult environment - that of the cart-pole environment of the OpenAI gym. In this environment the agent must balance a pole atop a cart. The pole follows normal laws of physics and generally topples rapidly down from its starting point where it points directly upwards. If the pole topples more than 14 degrees from the vertical the trial is a failure and the environment resets. The agent has two actions available - it can push the pole to the right or the left - and it must keep the pole upright for as long as possible. The input to the agent is four numbers encoding the position of the cart and the angle of the pole, and the first derivatives of the position and the angle.

Although this task is not particularly difficult for state of the art reinforcement learning agents, it poses a significantly greater challenge than the thermostat task. The agent must both learn from scratch a model of the environment, and also a model of how which action it takes affects the environment so that it compute which action to take at any point to minimize its variational free energy.

The preliminary cart-pole agent consists of a full construct model consisting of three hierarchical layers and two dynamical ones. The network is entirely linear. Additionally, the agent possesses a simple action model consisting of a single predictive linear layer which it can use to learn the effect of an action on the incoming sense data. It then selects

the action which will minimize the total free energy.

Training graphs of a single representative run of the cart-pole agent are shown below. The agent generally starts out chaotically but gradually learns to control the cart-pole until it reaches a baseline level of control which is not particularly good, but significantly better than random or simple reactive control (i.e. push left when pole is tipped to the right and vice-versa).

A video of the agent learning to act in the cart-pole environment can be found in the supplementary materials.

(a) The sense data given to the cartpole agent.

(b) The cartpole agent's predictions of the sense data

(c) The cartpole agent's action (0 or 1) at each time step

(d) The cartpole agents' prediction error at the lowest hierarchical level

(e) The cartpole's prediction error at the second hierarchical layer

(f) The cartpole agent's prediction error at the highest hierarchical layer
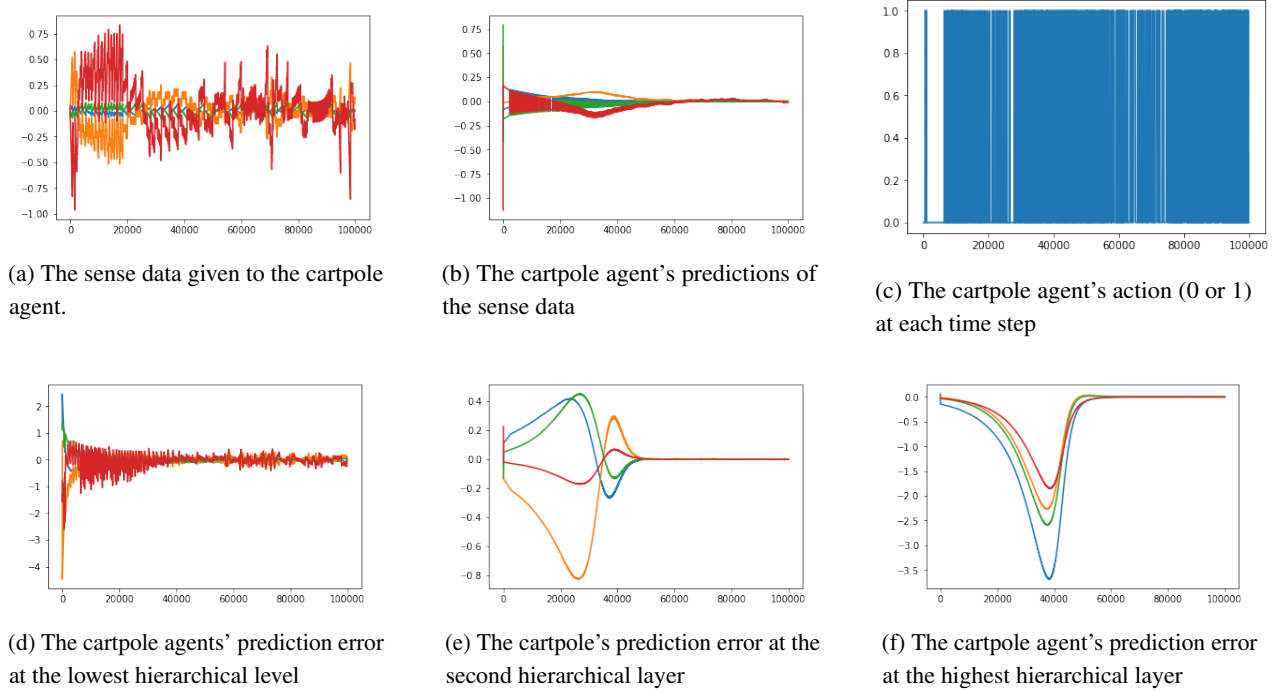
Figure 11: The training graphs of the active inference cart-pole model. The agent manages to make some headway successfully predicting and controlling the cart-pole. It is not perfect however, and significant prediction errors still remain. Much work has still to be done before an active inference agent can master the cart-pole environment.

There are several interesting things in the training graphs above. The first is that the agent does manage to learn and keep the cart-pole upright to some extent. It is definitely not perfect, however, and unlike the thermostat agent it does not appear to be asymptotically approaching the correct solution. The agent manages to successfully minimize the prediction errors in the top two levels of its hierarchy although significant prediction error is still maintained in the lower levels. This is not surprising as the precise way the pole topples is very difficult to predict in exacting detail. The relatively low maximum performance of the cart-pole agent could be due to the small scale and linearity of its perceptual model. The agent's "neural network" was only composed of 4 neurons for each hierarchical and dynamical layer - utterly tiny compared to typical deep learning approaches, and its linearity means that it is only able to predict linear trajectories into the future. These limitations likely significantly handicapped the performance of the agent, and offer much potential for scaling up in the future.

We have experimented with extending the baseline model in several ways. We have included nonlinear activation functions such as sigmoid, tanh, and relu functions but so far these tweaks appear to have had relatively little effect

on the performance of the agent. Further work is needed here to scale up the cartpole agent and then to extend and generalise the method to larger higher dimensional methods. It is possible that the current limitations of the agent are merely a reflection of a lack of size and thus neural network capacity. Additionally, the agent should be tested on other tasks to determine its performance in a wide range of environments. Further work on the cart-pole agent is forthcoming and a preprint can be found here: https://psyarxiv.com/kf6wc/.

Overall, therefore, we have implemented simple active inference agent that is able to successfully solve a toy task. We have also implemented a considerably more complex active inference agent comprising both a generative world model and an action model - both of which are predictive processing models of the type described in this paper. This agent is able to learn a reasonable, and significantly better than random policy, at a real, albeit simple, reinforcement learning task. It does this in an entirely unsupervised way and indeed without the rewards provided by the task environment. It is still limited and far from perfect on this task, both in terms of perceptual prediction and its choices of actions. Nevertheless, it shows real progress. Future work here will be focused on extending and scaling up the agent to enable it to solve the cart-pole task and then applying the architecture to progressively more demanding reinforcement learning tasks.

5 Conclusions and General Thoughts

We have implemented a variety of predictive processing models proposed in the literature, including static, dynamic, and "full-construct" predictive coding models, and also two different active inference agents. We have demonstrated that static predictive coding models can learn to reconstruct images, interpolate between them, and generate unseen digits via sampling their latent space. We have seen that dynamical predictive coding models can learn to predict sine waves and other simple wave forms, and we have implemented simple active inference agents with dynamical predictive coding perceptual models that are able to solve simple reinforcement learning tasks such as the cart-pole. These results provide empirical validation of much theoretical work in predictive coding by demonstrating that the models proposed in the literature can be implemented, do converge to minimise prediction error, and can be used to predict dynamic real-world stimuli. The challenge now will largely be one of scaling up and improving the stability of the models to enable them to attack more challenging task domains. Additionally, another challenge is trying to move beyond the baseline architectures presented here and to try to come up with new predictive processing architectures more carefully specialized for specific tasks such as sequence prediction or image recognition/generation. A further challenge is the implementation and scaling up of active inference agents to take on non-toy reinforcement learning tasks. We believe this is a promising avenue for further research since, firstly, a successful active inference agent requires a good predictive world model, so research here will also drive the development of better predictive perceptual models and secondly, predictive processing is a paradigm that fundamentally sees embodiment and interaction with a constantly changing world as a core process, and nowhere is that better instantiated than in an active inference agent that is constantly acting in and interacting with the world.

References

Baltieri, M. and Buckley, C. L. (2017). An active inference implementation of phototaxis. In *Artificial Life Conference Proceedings 14*, pages 36–43. MIT Press.

Bogacz, R. (2017). A tutorial on the free-energy framework for modelling perception and learning. *Journal of mathematical psychology*, 76:198–211.

Brown, H., Friston, K. J., and Bestmann, S. (2011). Active inference, attention, and motor preparation. *Frontiers in psychology*, 2:218.

Buckley, C. L., Kim, C. S., McGregor, S., and Seth, A. K. (2017). The free energy principle for action and perception: A mathematical review. *Journal of Mathematical Psychology*, 81:55–79.

Friston, K. (2003). Learning and inference in the brain. *Neural Networks*, 16(9):1325–1352.

Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 360(1456):815–836.

Friston, K. (2008a). Hierarchical models in the brain. *PLoS computational biology*, 4(11):e1000211.

Friston, K. (2009). The free-energy principle: a rough guide to the brain? *Trends in cognitive sciences*, 13(7):293–301.

Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature reviews neuroscience*, 11(2):127.

Friston, K. and Ao, P. (2012). Free energy, value, and attractors. *Computational and mathematical methods in medicine*, 2012.

Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., and Pezzulo, G. (2017). Active inference: a process theory. *Neural computation*, 29(1):1–49.

Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., Pezzulo, G., et al. (2016). Active inference and learning. *Neuroscience & Biobehavioral Reviews*, 68:862–879.

Friston, K., Mattout, J., and Kilner, J. (2011). Action understanding and active inference. *Biological cybernetics*, 104(1-2):137–160.

Friston, K., Mattout, J., Trujillo-Barreto, N., Ashburner, J., and Penny, W. (2007). Variational free energy and the laplace approximation. *Neuroimage*, 34(1):220–234.

Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., and Pezzulo, G. (2015). Active inference and epistemic value. *Cognitive neuroscience*, 6(4):187–214.

Friston, K., Schwartenbeck, P., FitzGerald, T., Moutoussis, M., Behrens, T., and Dolan, R. J. (2013). The anatomy of choice: active inference and agency. *Frontiers in human neuroscience*, 7:598.

Friston, K., Stephan, K., Li, B., and Daunizeau, J. (2010). Generalised filtering. *Mathematical Problems in Engineering*, 2010.

Friston, K. J. (2008b). Variational filtering. *NeuroImage*, 41(3):747–766.

Friston, K. J., Daunizeau, J., and Kiebel, S. J. (2009). Reinforcement learning or active inference? *PloS one*, 4(7):e6421.

Friston, K. J. and Stephan, K. E. (2007). Free-energy and the brain. *Synthese*, 159(3):417–458.

Karl, F. (2012). A free energy principle for biological systems. *Entropy*, 14(11):2100–2121.

McGregor, S., Baltieri, M., and Buckley, C. L. (2015). A minimal active inference agent. *arXiv preprint arXiv:1503.04187*.

Pezzulo, G., Rigoli, F., and Friston, K. (2015). Active inference, homeostatic regulation and adaptive behavioural control. *Progress in neurobiology*, 134:17–35.

Rao, R. P. and Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience*, 2(1):79.

Spratling, M. W. (2008). Reconciling predictive coding and biased competition models of cortical function. *Frontiers in computational neuroscience*, 2:4.

Spratling, M. W. (2017). A hierarchical predictive coding model of object recognition in natural images. *Cognitive computation*, 9(2):151–167.